

Lecture 19: Encoding problems and decidability

7 April 2009

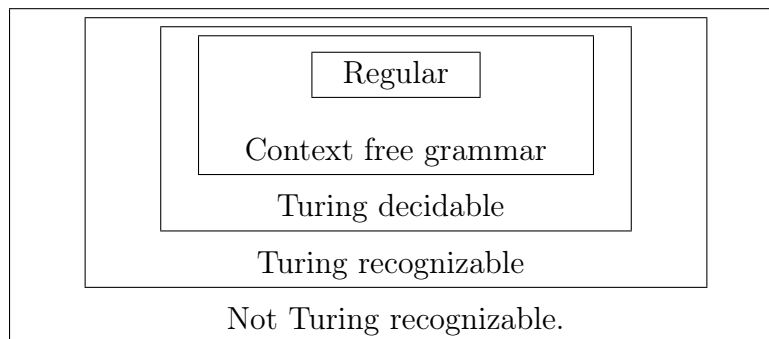
This lecture presents examples of languages that are *Turing decidable*.

1 Review and context

Remember that a Turing machine M can do three sorts of things on an input w . The TM M might halt and accept. It might halt and reject. Or it might never halt. A TM is a *decider* if it always halts on all inputs.

A TM *recognizable* language is a language L for which there is a TM M , such that $L(M) = L$. A TM decidable language is a language L for which there is a decider TM M , such that $L(M) = L$.

Here is a figure showing the hierarchy of languages.



Conceptually, when we think about algorithms in computer science, we are normally interested in code which is guaranteed to halt on all inputs. So, for questions about languages, our primary interest is in Turing decidable (not just recognizable) languages.

Any algorithmic task can be converted into decision problem about languages. Some tasks are naturally in this form, e.g. “Is the length of this string prime?”. In other cases, we have to restructure the question in one of two ways:

- A complex input object (e.g. a graph) may need to be encoded as a string.
- A construction task may have to be rephrased as a yes/no question.

2 TM example: Adding two numbers

2.1 A simple decision problem

For example, consider the task of adding two decimal numbers. The obvious algorithm might take two numbers a and b as input, and produce a number c as output. We can rephrase this as a question about languages by asking “Given inputs a , b , and c , is $c = a + b$ ”.

For the alphabet

$$\Sigma = \{0, 1, \dots, 9, +, -\},$$

consider the language

$$L = \left\{ a_n a_{n-1} \dots a_0 + b_m b_{m-1} \dots b_0 = c_r c_{r-1} \dots c_0 \left| \begin{array}{l} a_i, b_j, c_k \in [0, 9] \text{ and} \\ \langle a_n a_{n-1} \dots a_0 \rangle \\ + \quad \langle b_m b_{m-1} \dots b_0 \rangle \\ = \langle c_r c_{r-1} \dots c_0 \rangle \end{array} \right. \right\},$$

where $\langle a_n a_{n-1} \dots a_0 \rangle = \sum_{i=0}^n a_i \cdot 10^i$ is the number represented in base ten by the string $a_n a_{n-1} \dots a_0$.

We then ask whether we can build a TM which decides the language L .

2.2 A decider for addition

To build a decider for this addition problem, we will use a multi-tape TM. We showed (last class) that a multi-tape TM is equivalent to a single tape TM. First, let us build a useful helper function, which reverses the contents of one tape.

2.2.1 Reversing a tape

Given the content of tape $\textcircled{1}_1$, we can reverse it easily in two steps using a temporary tape. First, we put a marker onto the temporary tape. Moving the heads on both tapes to the right, we copy the contents of $\textcircled{1}_1$ onto the temporary tape.

Next, we put the $\textcircled{1}_1$ head at the start of its tape, but the temporary tape head remains at the end of this tape. We copy the material back onto $\textcircled{1}_1$, but in reverse order, moving the $\textcircled{1}_1$ head rightwards and the temporary tape head leftwards.

Let **ReverseTape**(t) denote the TM mechanism (i.e. procedure) that reverses the t th tape. We are going to buildup TM by putting together such procedures.

2.2.2 Adding two numbers

Now, let us assemble the addition algorithm. We will use five tapes: the input ($\textcircled{1}_1$), three tapes to hold numbers ($\textcircled{2}_2$, $\textcircled{3}_3$, and $\textcircled{4}_4$), and a scratch tape used for the reversal operation.

The TM will first scan the input tape (i.e., $\textcircled{1}_1$), and copy the first number to $\textcircled{2}_2$, and the second number to $\textcircled{3}_3$. Next, we do **ReverseTape**(2) and **ReverseTape**(3). Now, we move the head of $\textcircled{2}_2$ and $\textcircled{3}_3$ to the beginning of the tapes, and we start moving them together computing the sum of the digits under the two heads, writing the output to $\textcircled{4}_4$, and moving the three heads to the right. Naturally, we have a carry over digit, which we encode in the current state of the TM controller (the carry over digit is either 0, 1 or 2).

If one of the heads of $\textcircled{2}_2$ or $\textcircled{3}_3$ reaches the end of the tape, then we continue moving it, interpreting $_$ as a 0. We halt when the heads on both tapes see $_$.

Next, we move the head of $\textcircled{4}_4$ back to the beginning of the tape, and do **ReverseTape**(4). Finally, we compare the content of $\textcircled{4}_4$ with the number written on $\textcircled{1}_1$ after the = character. If they are equal, the TM accepts, otherwise it rejects.

3 Encoding a graph problem

As the above example demonstrates, the *coding scheme* used for the input has big impact on the complexity of our algorithm. The addition algorithm would have been easier if the numbers were written in reverse order, or if they had been in binary. Such details may affect the running time of the algorithm, but they do not change whether the problem is Turing decidable or not.

When algorithms operate on objects that are not strings, these objects need to be *encoded* into strings before we can make the algorithm into a decision problem. For example, consider the following situation. We are given a *directed graph* $G = (V, E)$, and two vertices $s, t \in V$, and we would like to decide if there is a way to reach t from s .

All sorts of encodings are possible. But it is easiest to understand if we use encodings that look like standard ASCII file, of the sort you might use as input to your Java or C++ program. ASCII files look like they are two-dimensional. But remember that they are actually one-dimensional strings inside the computer. Line breaks display in a special way, but they are underlyingly just a special separator character (`<NL>` on a unix system), very similar to the `$` or `#` that we've used to subdivide items in our string examples.

To make things easy, we will number the vertices of V from 1 to $n = |V|$. To specify that there is an edge between two vertices u and v , we then specify the two indices of u and v . We will use the notation (u, v) . Thus, to specify a graph as a text file, we could use the following format, where n is the number of vertices and m is the number of edges in the graph.

```
n
m
(n1, n'1)
(n2, n'2)
⋮
(nm, n'm)
```

Namely, the first line of the file, will contain the number (written explicitly using ASCII), next the second line is the number of edges of G (i.e., m). Then, every line specify one edge of the graph, by specifying the two numbers that are the vertices of the edge. As a concrete example, consider the following graph.

The number of edges is a bit redundant, because we could just stop reading at the end of the file. But it is convenient for algorithm design.

See Figure 1, for an example of a graph its encoding using these scheme.

4 Algorithm for graph reachability

To encode an instance of the *s, t-reachability problem*, our ASCII file will need to contain not only the graph but also the vertices s and t . The input tape for our TM would contain all this information, laid out in 1D (i.e. imagine the line break displayed as an ordinary separator character).




Figure 1: A graph encoded as text. The string encoding the graph is in fact “5⟨NL⟩7⟨NL⟩(1,2)⟨NL⟩(2,3)⟨NL⟩(3,5)⟨NL⟩(5,1)⟨NL⟩(3,4)⟨NL⟩(4,3)⟨NL⟩(4,2)”. Here ⟨NL⟩ denotes the special new-line character.

To solve this problem, we will need to search the graph, starting with node s . The TM accepts iff this search finds the node t . We will store information on four TM tapes, in addition to the input tape. The TM would have the following tapes:

- ⊗₁: Input tape
- ⊗₂: Target node t .
- ⊗₃: Edge list.
- ⊗₄: *Done list*: list of nodes that we’ve finished processing
- ⊗₅: *To-do list*: list of nodes whose outgoing edges have not been followed

Given the graph, the TM reads the graph (checking that the input is in the right format). It puts the list of edges onto tape ⊗₃, puts t onto its own tape (i.e., ⊗₂), and puts the node s onto the to-do list tape (i.e., ⊗₅).

Next, the TM loops. In each iteration, it removes the first node x from the to-do list. If $x = t$, the TM halts and accepts. Otherwise, x is added to the done list (i.e., ⊗₄). Then the TM searches the Edge list for all edges going outwards from x . Suppose an outgoing edge goes from x to y . Then if y is not already on the finished list or the to-do list, then y is added to the to-do list.

If there is nothing left on on the to-do list, the TM halts and rejects.

This algorithm is a **graph search** algorithm. It is breadth-first search if the new nodes are added to the end of the to-do list and depth-first search if they are added in the start of the list. (Or, said another way, the to-do list operates as either a queue or a stack.)

The separate visited list is necessary to prevent the algorithm from going into an infinite loop if the graph contains cycles.

5 Some decidable DFA problems

If M is a DFA, the string encoding of M is written as $\langle M \rangle$.

The string encoding of a DFA is similar to the encoding of a directed graph except that our encoding has to have a label for each edge, specify the start state, and list the final states.

Emptiness of DFA. Consider the language

$$E_{\text{DFA}} = \left\{ \langle M \rangle \mid M \text{ is a DFA, and } L(M) = \emptyset \right\}.$$

This language is decidable. Namely, given an instance $\langle M \rangle$, there is a TM that reads $\langle M \rangle$, this TM always stops, and accepts if and only if $L(M)$ is empty. Indeed, do a graph search on the DFA (as above) starting at the start state of M , and check whether any of the final states is reachable. If so, the $L(M) \neq \emptyset$.

Lemma 5.1 *The language E_{DFA} is decidable.*

Emptiness of NFA. Consider the following language

$$E_{\text{NFA}} = \left\{ \langle M \rangle \mid M \text{ is a NFA, and } L(M) = \emptyset \right\}.$$

This language is decidable. Indeed, convert the given NFA into a DFA (as done in class, long time ago) and then call the code for E_{DFA} on the encoded DFA. Notice that the first step in this algorithm takes the encoded version of M and writes the encoding for the corresponding DFA. You can imagine this as taking a state diagram as input and producing a new state diagram as output.

Equal languages for DFAs. Consider the language

$$EQ_{\text{DFA}} = \left\{ \langle M, D \rangle \mid M \text{ and } D \text{ are NFAs, and } L(M) = L(D) \right\}.$$

This language is also decidable. Remember that the *symmetric difference* of two sets X and Y is $X \oplus Y = (X \cap \bar{Y}) \cup (Y \cap \bar{X})$. The set $X \oplus Y$ is empty if and only if the two sets are equal. But, given a DFA, we know how to make a DFA recognizing the complement of its language. And we also know how to take two DFA's and make a DFA recognizing the union or intersection of their languages.

So, given the encodings for M and D , our TM will construct the encoding of a DFA $\langle N \rangle$ recognizing the symmetric difference of their languages. Then it would call the code for deciding if $\langle N \rangle \in E_{\text{DFA}}$.

Informally, problems involving regular languages are always decidable, because they are so easy to manipulate. Problems involving context-free languages are sometimes decidable. And only the simplest problems involving Turing machines are decidable.

6 The acceptance problem for DFA's

The following language is also decidable:

$$A_{\text{DFA}} = \left\{ \langle M, w \rangle \mid M \text{ is a DFA, } w \text{ is a word, and } M \text{ accepts } w. \right\}.$$

As before, the notation $\langle M, w \rangle$ is the encoding of the DFA M and the word w ; that is, it is the pair $\langle M \rangle$ and $\langle w \rangle$. For example, if $\langle w \rangle$ is just w (it's already a string), then $\langle M, w \rangle$ might be $\langle M \rangle \# w$ where $\#$ is some separator character. Or it might be $(\langle M \rangle, w)$. Or anything similar that encodes the input well. We will just assume that it is in some such reasonable encoding of a pair and that the low-level code for our TM (which we will not spell out in detail) knows what it is.

A Turing machine deciding A_{DFA} needs to be able to take the code for some arbitrary DFA, plus some arbitrary string, and decide if that DFA accepts that string. So it will need to contain a general-purpose DFA simulator. This is called the *acceptance problem* for DFA's.

It's useful to contrast this with a similar-sounding claim. If M is any DFA, then $L(M)$ is Turing-decidable. Indeed, to build a TM that accepts $L(M)$, we simply move the TM head to the right over the input, using the TM's controller to simulate the controller of the DFA directly.

In this case, we are given a specific fixed DFA M and we only need to cook up a TM that recognizes strings from this one particular language. This is much easier than A_{DFA} .

To decide A_{DFA} , our TM will use five tapes:

- ⊙₁: input: $\langle M, w \rangle$,
- ⊙₂: state,
- ⊙₃: final states
- ⊙₄: transition triples
- ⊙₅: input string.

The simulator then runs as follows:

- (1) Check the format of the input. Copy the start state to tape ⊙₂. Copy the input string to tape ⊙₅. Copy the transition triples and final states of the input machine $\langle M \rangle$ to tapes ⊙₃ and ⊙₄.
- (2) Put the tape ⊙₅ head at the beginning of the tape.
- (3) Find a transition triple $p \xrightarrow{c} q$ (written on tape ⊙₄) whose input state and character match the state written on tape ⊙₁ (i.e., p) and the character (i.e., c) under the head on tape ⊙₅.
- (4) Change the current state of the simulated DFA from p to q .

Specifically, copy the state q (written on the triple we just found on ⊙₄), to tape ⊙₂.

- (5) Move the tape \odot_5 head to the right (i.e., the simulation handled this input character).
- (6) Goto step (3).
- (7) Halt the loop when the tape \odot_5 head sees a blank. Accept if and only if the state on tape \odot_2 is one of the states on list of final states, stored on tape \odot_3 .